



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: Foster D. Hinshaw, David L. Meyers and Barry M. Zane
Application No.: 10/668,113 Group Art Unit: 2141
Filed: September 18, 2003 Examiner: Not assigned
Confirmation No.: 7658
Title: PROGRAMMABLE STREAMING DATA PROCESSOR FOR
DATABASE APPLIANCE HAVING MULTIPLE PROCESSING UNIT
GROUPS

CERTIFICATE OF MAILING OR TRANSMISSION	
I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as First Class Mail in an envelope addressed to Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, or is being facsimile transmitted to the United States Patent and Trademark Office on:	
<u>3/3/04</u>	<u>Nadene Z. Kush</u>
Date	Signature
<u>Nadene Z. Kush</u>	
Typed or printed name of person signing certificate	

PETITION TO MAKE SPECIAL FOR NEW APPLICATION

UNDER M.P.E.P. § 708.02, VII

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

Applicants hereby petition to make this new application, which has not received any examination by an Examiner, special as permitted under 37 CFR § 1.102(d). The Petition fee in the amount of \$130 is included in the enclosed check. Please charge any deficiency in fees and credit any overpayment to Deposit Account 08-0380. A duplicate copy of this Petition is enclosed for accounting purposes.

Applicants believe that all the claims in this application are directed to a single invention. If the Office determines that all the claims presented are not obviously directed to a single invention, then applicants will make an election without traverse as a prerequisite to the grant of special status.

A pre-examination search has been made by the inventors and their attorneys. The search strategy included reviewing publications known to the inventors and additional searching using terms such as “database”, “distributed computing”, “structured query language processing”, “parallel processing”, “streaming” and “tuple sets”.

A copy of each of the references deemed most closely related to the subject matter encompassed by the claims is provided. A Supplemental Information Disclosure Statement and accompanying PTO-1449 form are being filed concurrently with this Petition To Make Special.

Below is a detailed discussion of the present invention and the prior art references, which particularly points out how the claimed subject matter is distinguishable.

The present invention is a circuit that is capable of processing data from a streaming data source, such as a disk drive, prior to its being forwarded to a Central Processing Unit (CPU) and a more general purpose processor. Such processor that might be adapted for accepting queries that are requests for data in a database. Queries are typically provided in a Structured Query Language (SQL). In return a host processor develops a plan for executing the query, typically dividing the plan into a set of jobs to be executed by a number of distributed processing units called Job Processing Units (JPUs) in the application.

Each JPU has a special purpose programmable processor referred to as a Programmable Streaming Data Processor (PSDP) and also may include a general purpose local Central

Processing Unit (CPU) as well. The PSDP is programmable by the host and/or local CPU to interpret data in a specific format as it is read from the associated disks. This enables the PSDP to perform portions of jobs directly on data immediately, as it is read off the disk prior to such data ever being forwarded to the CPU for further processing.

For example, the PSDP can parse non-field delineated, streaming data from a mass storage device into block header fields, record header fields, and record data fields. The PSDP can also actually filter the record header and data field so that only certain fields on certain records are actually output by the PSDP to be placed in the JPU's memory. The PSDP thus can relieve the JPU from initial processing steps required to carry out an SQL query.

U.S. Patent 5,999,937 issued to Ellard discloses a customizable system for transferring data between an input data set and an output data set having possibly different data formats. The system allows conversion of a single input data record into a plurality of input data records to a single output data record. However, the system presumes that the records are already delineated. That is, there is no processing of non-field delineated data. Specifically, the field processing database 62 contains instructions for converting one field format to another but is not capable of processing non field delineated data as it streams from a storage medium.

U. S. Patent 5,884,299 to Ramesh *et al.* discloses a technique for optimizing SQL queries in a relational database management system using aggregate and grouping functions. A local aggregation operation is performed on one or more processors of a Massively Parallel Processor (MPP) computer system, wherein rows of the table that are local to each processor are locally aggregated to create one or more aggregate result rows. The aggregate result rows created by each of the local aggregation operations are then re-distributed to one or more processors, and coalesced into a global aggregate result row by a global aggregation operation.

Ramesh's optimization techniques involving involve expressions e.g., SUM, COUNT, MIN, MAX, AVG) and grouping (e.g., GROUP BY) clauses in SQL queries. He shows a Massively Parallel Processor (MPP) computer system which includes Data Storage Units (DSUs) and nodes. Each node includes one or more "virtual processors" such as Parsing Engines (PEs) and Virtual Access Module Processors (VAMPs). Each DSU stores some of the rows of each table. Using a hash function, the data storage can be evenly divided among the VAMPs. The PEs fully parallelize all functions among the VAMPs.

Ramesh does disclose parallel processing of local and global aggregation results, he does not teach the present invention, since his parser engines are not able to process non-field delineated data into record fields as it streams off a mass storage device.

U. S. Patent 5,937,401 issued to Hillegas describes a client server database system which processes an ordered tuple stream. Specifically, the system provides filtering to eliminate duplicate records without having to first perform a sort operation. The filter is implemented at the level of a query processor. It is shown in the diagram of Fig. 2 that this processor may be a database server 230 that is separate from the client processor 210. In particular, the database system 240 may be a Sybase SQL server type system available from Sybase, Inc. of Emeryville, California, which generally operates as an independent processor running under a server operating system, such as Microsoft Windows NT. An engine 260 within the database server system 240 includes a parser 261, normalizer 263, compiler 265, execution unit 268 and access methods 269.

The Hillegas system can thus operate as a filter to eliminate duplicate records without first having to perform a sort. The filter is only implemented at the level of the query processor (i.e., the filter is implemented as part of the software in the database server system 240), and not in hardware.

While the tables themselves are described as comprising horizontal rows or records (tuples) together with vertical columns or fields, there is no discussion of a separate programmable pipeline processor for processing streaming input data such as may be received from a mass storage device, parsing non-field delineated data from the streaming source into field delineated data, and then an additional field buffer to store the extracted fields and/or logic units to perform field operations on the extracted fields.

U. S. Patent 5,937,415 issued to Sheffield, *et al.* discloses a client server database system for providing improved methods for executing database queries. A pipeline object facilitates moving data from one database management system to another or between databases of the same type. More specifically, as shown in Fig. 2, the database server system 240 is a Sybase SQL server type system that is generally implemented as an independent software process running under a server operated system such as Microsoft Windows NT. In this arrangement, clients 210 store data in or retrieve data from one or more database tables 250 under control of the database server system 240. The Sheffield system is almost identical to the Hillegas system, is implemented entirely in software, and does not teach the benefits of distributed hardware functionality.

U. S. Patent 6,138,118 issued to Koppstein, *et al.* discloses a system for reconciling the execution of a high priority stream of instructions that is concurrently executed with a low priority stream of transactions. In particular, a scheduling database generally comprises resources, tasks, and associated assignment relations, all which are governed by various timing rules, for example, by specific time windows for completing certain tasks.

U. S. Patent 6,339,772 issued to Klein, *et al.* discloses an SQL compiler and SQL executor for a database management system that can be extended to process queries requiring streaming or processing of data stored on a table. According to this patent, an SQL compiler and

SQL executor are extended to process operations on streams of tuples and to access regular database tables as continuous streams of tuples. In particular, the SQL executor first reads all qualifying tuples in a specified table, and subsequently monitors and returns new qualifying tuples. A first part of this method is performed by a regular table scan, while the second part of the method is performed by a delta scan. The monitoring function is performed until the cursor representing the SQL statement being executed, including the scan operations, is closed by the calling application. Execution of the SQL statement suspends and automatically resumes (i.e., as a reschedule) when new data becomes available. As shown in the drawings, particularly at Fig. 3, partitions of the database 58 may be stored on different nodes of a relational database system. The application process 60 represents a process or processes that execute not only the application program but also portions of the execution tree 54 above the leaf nodes. The leaf nodes are executed by disk processes 62 in each of the nodes of the transaction processing system.

U. S. Patent 6,415,373 issued to Peters, *et al.* discloses a data processing system in which multiple high bandwidth streams of data can be distributed among multiple storage units.

U. S. Patent 6,493,701 issued to Ponnekanti discloses a database system for providing nested loop “join” operations. The system is implemented in a database server system software process. An SQL query from a client database application can specify a join of three or more tables. Where at least one join condition exists between an inner table and an outer table, and is not immediately or directly preceding table, the joining order itself specifies the particular order or sequence that the tables are accessed for retrieving rows for examination during query execution.

In particular, a loop is established to retrieve rows from successive tables (specified by the join order). The method then determines whether a condition is being tested that refers back to an more outer table that is not a directly preceding table. If this condition is not met, then

query execution proceeds to fetch the next row from that outer table. Otherwise, the method continues down the join order to examine any remaining or subsequent tables in the join order, applying any subsequent query conditions that must be met in order to qualify for the query.

U. S. Patent 6,507,834 issued to Kabra, *et al.* discloses a technique for parallel execution of SQL operations from stored procedures. The query is optimized, parallelized, and then processed by a dispatcher. The dispatcher executes the query by setting up communication links between various operators in the query, insuring that results are sent back to the data server that originated the query request. The dispatcher can merge the results of parallel execution producing a single stream of tuples, for example, that is returned to the calling stored procedure. Queries are carried out on two classes of processes including query controllers 104 and data servers 130.

U. S. Patent 6,542,508 issued to Lin describes a hardware based policy engine and employs a policy cache to process packets of network traffic. The policy engine includes a stream classifier that associates each packet with at least one action processor based on data in the packet.

U. S. Patent 6,542,886 issued to Chadhuri, *et al.* discloses a database server system that supports weighted and unweighted sampling of records or tuples in accordance with sampling semantics such as sample with replacement, without replacement, or independent coin flips sampling semantics. A server may perform such sampling sequentially not only to sample non-materialized records such as those produced as a stream in a pipeline in a query tree, for example, but also to sample records whether materialized or not in a single pass.

U. S. Patent Publication US 2002/0038313 by Klein, *et al.* also proposes an SQL compiler and SQL executor in a database management system that is extended to process queries

requiring a streaming mode type of processing of data stored in a table. As with the previously described U.S. Patent 6,339,772, a scan operator performs initial scans to access rows in a specified database table and then performs a delta scan to access new rows added to the table as well as rows modified by other queries.

U. S. Patent Publication US 2002/0161748 by Hamel, *et al.* discloses a method for loading data from a remote data source on a record by record basis. This system uses structured query language statements to request data loaded from a remote source. Data are transported, record by record, via a database connection communication line, where a target site loads records concurrently with unloading of records in the source site. The data loading is thus performed in pipeline manner, with the plurality of parallel streams pointed to by a plurality of data source partition cursors.

The following references are only of more general interest.

U.S. Patent No. 4,594,655 issued to Hao *et al.* describes a secondary data flow facility that emulates certain operations in an instruction pipeline.

U.S. Patent 4,811,214 issued to Nosenchuck *et al.* discloses a multimode parallel processing computer system wherein each node is a large capacity node that includes a reconfigurable pipeline of functional units.

U.S. Patent Publication 2001/0036322 issued to Bloomfield *et al.* discloses an image processing system that is not concerned with processing fields and/or records in databases.

U.S. Patent Publication 2002/0095400 issued to Johnson *et al.* discloses techniques for delivering services in a network environment.

U.S. Patent Publication 2002/0128823 issued to Kovacevic discloses a system for processing digital audio streams received from transport streams. An audio parser discards transport packets not related to specific audio types providing packetized elementary stream audio data to an audio decoding system. As with other prior art, this patent is not concerned with processing non field delineated data as it streams from a mass storage device.

In summary, none of the above patents disclose the system, apparatus, method or computer program product of the present invention. We therefore ask that this Petition to Make Special for the present application be granted.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 

David J. Thibodeau, Jr.
Registration No. 31,671
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: 3/3/04